



LEGAL NOTICE:

© Copyright 2007 - 2015 NVM Express, Inc. ALL RIGHTS RESERVED.

This NVM Express revision 1.2 specification is proprietary to the NVM Express, Inc. (also referred to as "Company") and/or its successors and assigns.

NOTICE TO USERS WHO ARE NVM EXPRESS, INC. MEMBERS: Members of NVM Express, Inc. have the right to use and implement this NVM Express revision 1.2 specification subject, however, to the Member's continued compliance with the Company's Intellectual Property Policy and Bylaws and the Member's Participation Agreement.

NOTICE TO NON-MEMBERS OF NVM EXPRESS, INC.: If you are not a Member of NVM Express, Inc. and you have obtained a copy of this document, you only have a right to review this document or make reference to or cite this document. Any such references or citations to this document must acknowledge NVM Express, Inc. copyright ownership of this document. The proper copyright citation or reference is as follows: "**© 2007 - 2015 NVM Express, Inc. ALL RIGHTS RESERVED.**" When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

LEGAL DISCLAIMER:

THIS DOCUMENT AND THE INFORMATION CONTAINED HEREIN IS PROVIDED ON AN "**AS IS**" BASIS. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, NVM EXPRESS, INC. (ALONG WITH THE CONTRIBUTORS TO THIS DOCUMENT) HEREBY DISCLAIM ALL REPRESENTATIONS, WARRANTIES AND/OR COVENANTS, EITHER EXPRESS OR IMPLIED, STATUTORY OR AT COMMON LAW, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, VALIDITY, AND/OR NONINFRINGEMENT.

All product names, trademarks, registered trademarks, and/or servicemarks may be claimed as the property of their respective owners.

NVM Express Workgroup
c/o Virtual, Inc.
401 Edgewater Place, Suite 600
Wakefield, MA 01880
info@nvmexpress.org

NVM Express™ Technical Errata

Errata ID	002
Revision Date	3/7/2015
Affected Spec Ver.	NVM Express™ 1.2
Corrected Spec Ver.	

Errata Author(s)

Name	Company
John Carroll	Intel
Dave Landsman	SanDisk
Koichi Nagai	Toshiba
Fred Au	Toshiba
Jim Hatfield	Seagate
Ramdas Kachare	Seagate

Errata Overview

A recommendation is made for the host to avoid issuing I/O when selecting a non-operational power state to avoid unexpected results.

Several editorial changes are included, as well as reference fixes.

A clarification is made that a Controller Reset does not affect the MMIO MSI-X register state.

A clarification is made that RPMB targets are not part of namespaces.

A clarification is made that for Dataset Management, all command attribute combinations may be set.

A clarification is made that Firmware Activation Notices are sent to any controller affected by a Firmware Commit that is activated immediately.

Revision History

Revision Date	Change Description
11/18/2014	First draft.
12/18/2014	Added firmware activate notice change and Figure 11 update based on WG feedback
1/8/2015	Added Host memory buffer change, updated based on 1/8 WG feedback
1/14/2015	Minor editorial changes.
3/7/2015	Ratified.

Description of Specification Changes

Modify a portion of section 5.14.1 as shown below:

5.14.1 Feature Specific Information

Figure 108 defines the Features that may be configured with Set Features and retrieved with Get Features. Figure 109 defines Features that are specific to the NVM Command Set. Some Features utilize a memory buffer to configure or return attributes for a Feature, whereas others only utilize a Dword in the command or completion queue entry. Feature values that are not persistent across power states **and resets** are reset to their default values as part of a controller reset operation.

Modify a portion of section 8.4.1 as shown below:

8.4.1 Non-Operational Power States

A power state may be a non-operational power state, as indicated by Non-Operational State (NOPS) field in Figure **90 91**. In a non-operational power state, memory-mapped I/O (MMIO) accesses, configuration register accesses and Admin Queue commands are serviced. No I/O commands are processed by the controller while in a non-operational power state. **The host should wait until there are no pending I/O commands prior to issuing a Set Features command to change the current power state of the device to a non-operational power state, and not submit new I/O commands until the Set Features command completes. Issuing an I/O command in parallel may result in the controller being in an unexpected power state.**

Modify a portion of section 2.2 as shown below:

2.2 PCI Power Management Capabilities

See **section 3.1.5** for requirements when the PCI power management state changes.

Modify a portion of Figure 55 as shown below:

15:00	Queue Identifier (QID): This field indicates the identifier to assign to the Submission Queue to be created. This identifier corresponds to the Submission Queue Tail Doorbell used for this command (i.e., the value y in section 0 section 3.1.14). This value shall not exceed the value reported in the Number of Queues feature (see section Error! Reference source not found.) for I/O Submission Queues. If the value specified is 0h, exceeds the Number of Queues reported, or corresponds to an identifier already in use, the controller should return an error of Invalid Queue Identifier.
-------	--

Modify a portion of section 3.1.5 as shown below:

00	RW	0	<p>Enable (EN): When set to '1', then the controller shall process commands based on Submission Queue Tail doorbell writes. When cleared to '0', then the controller shall not process commands nor post completion queue entries to Completion Queues. When this field transitions from '1' to '0', the controller is reset (referred to as a Controller Reset). The reset deletes all I/O Submission Queues and I/O Completion Queues, resets the Admin Submission Queue and Completion Queue, and brings the hardware to an idle state. The reset does not affect PCI Express registers (including MMIO MSI-X registers), nor the Admin Queue registers (AQA, ASQ, or ACQ). All other controller registers defined in this section and internal controller state (e.g., Feature values defined in section 0 5.14.1 that are not persistent across power states) are reset to their default values. The controller shall ensure that there is no data loss for commands that have had corresponding completion queue entries posted to an I/O Completion Queue prior to the reset operation. Refer to section Error! Reference source not found. for reset details.</p> <p>When this field is cleared to '0', the CSTS.RDY bit is cleared to '0' by the controller once the controller is ready to be re-enabled. When this field is set to '1', the controller sets CSTS.RDY to '1' when it is ready to process commands. CSTS.RDY may be set to '1' before namespace(s) are ready to be accessed.</p> <p>Setting this field from a '0' to a '1' when CSTS.RDY is a '1,' or setting this field from a '1' to a '0' when CSTS.RDY is a '0,' has undefined results. The Admin Queue registers (AQA, ASQ, and ACQ) shall only be modified when EN is cleared to '0'.</p>
----	----	---	--

Modify a portion of section 8.10 as shown below:

The controller may support multiple RPMB targets. **RPMB targets are not contained within a namespace. Security Send and Security Receive commands for RPMB do not use the namespace ID field; NSID shall be cleared to 0h.** Each RPMB target operates independently – there may be requests outstanding to multiple RPMB targets at once (where the requests may be interleaved between RPMB targets). In order to guarantee ordering the host should issue and wait for completion for one Security Send or Security Receive command at a time. Each RPMB target requires individual authentication and key programming. Each RPMB target may have its own unique Authentication Key.

Modify a portion of section 6.7 as shown below:

Figure 1: Dataset Management – Command Dword 11

Bit	Description
31:03	Reserved

02	Attribute – Deallocate (AD): If set to '1' then the NVM subsystem may deallocate all provided ranges. If a read occurs to a deallocated range, the controller shall return all zeros, all ones, or the last data written to the associated LBA. If the deallocated or unwritten logical block error is enabled and a read occurs to a deallocated range, then the read shall fail with the Unwritten or Deallocated Logical Block status code.
01	Attribute – Integral Dataset for Write (IDW): If set to '1' then the dataset should be optimized for write access as an integral unit. The host expects to perform operations on all ranges provided as an integral unit for writes, indicating that if a portion of the dataset is written it is expected that all of the ranges in the dataset are going to be written.
00	Attribute – Integral Dataset for Read (IDR): If set to '1' then the dataset should be optimized for read access as an integral unit. The host expects to perform operations on all ranges provided as an integral unit for reads, indicating that if a portion of the dataset is read it is expected that all of the ranges in the dataset are going to be read.

If the Dataset Management command is supported, all combinations of attributes specified in Figure 166 may be set.

The data that the Dataset Management command provides is a list of ranges with context attributes.

Modify a portion of section 5.14.1.3 as shown below:

Each entry in the LBA Range Type data structure is defined in **Error! Reference source not found.**. The LBA Range feature is a set of 64 byte entries; the number of entries is indicated as a command parameter, the maximum number of entries is 64. The LBA ranges shall not overlap. If the LBA ranges overlap, the controller should return an error of Overlapping Range. All unused entries in the LBA Range Type data structure shall be cleared to all zeroes for both Get Features and Set Features.

The default value for this feature should clear the Number of LBA Ranges field to 00h and initialize the LBA Range Type data structure to contain a single entry with:

- Type field cleared to 00h,
- Attributes field set to 01h,
- Starting LBA field cleared to 0h,
- Number of Logical Blocks field set to indicate the number of LBAs in the namespace, and
- GUID field set to a globally unique identifier.

Modify a portion of section 8.1 as shown below:

The host submits a Firmware Commit command with a Commit Action of 011b which specifies that the image should be activated immediately without reset. If the controller starts to activate the firmware ~~and Firmware Activation Notices are enabled (refer to Figure 122), any the controllers affected by the new firmware send~~ a Firmware Activation Starting asynchronous event to the host ~~if Firmware Activation Notices are enabled (refer to Figure 122).~~

Modify Figure 11 as shown below (resort the lines so the lowest byte number is at the top):

Figure 11: Command Format – Admin Command Set

Bytes	Description
03:00	Command Dword 0 (CDW0): This field is common to all commands and is defined in Figure 10.
07:04	Namespace Identifier (NSID): This field specifies the namespace ID that this command applies to. If the namespace ID is not used for the command, then this field shall be cleared to 0h. If a command shall be applied to all namespaces accessible by this controller, then this field shall be set to FFFFFFFFh. Unless otherwise noted, specifying an inactive namespace ID in a command that uses the namespace ID shall cause the controller to abort the command with status Invalid Field in Command. Specifying an invalid namespace ID in a command that uses the namespace ID shall cause the controller to abort the command with status Invalid Namespace or Format.
15:08	Reserved
23:16	Metadata Pointer (MPTR): This field contains the address of a contiguous physical buffer of metadata. This field is only used if metadata is not interleaved with the logical block data, as specified in the Format NVM command. This field shall be Dword aligned.
31:24	PRP Entry 1 (PRP1): This field contains the first PRP entry for the command or a PRP List pointer depending on the command.
39:32	PRP Entry 2 (PRP2): This field: a) is reserved if the data transfer does not cross a memory page boundary. b) specifies the Page Base Address of the second memory page if the data transfer crosses exactly one memory page boundary. E.g.,: i. the command data transfer length is equal in size to one memory page and the offset portion of the PBAO field of PRP1 is non-zero or ii. the Offset portion of the PBAO field of PRP1 is equal to zero and the command data transfer length is greater than one memory page and less than or equal to two memory pages in size. c) is a PRP List pointer if the data transfer crosses more than one memory page boundary. E.g.,: i. the command data transfer length is greater than or equal to two memory pages in size but the offset portion of the PBAO field of PRP1 is non-zero or ii. the command data transfer length is equal in size to more than two memory pages and the Offset portion of the PBAO field of PRP1 is equal to zero.
43:40	Command Dword 10 (CDW10): This field is command specific Dword 10.
47:44	Command Dword 11 (CDW11): This field is command specific Dword 11.
51:48	Command Dword 12 (CDW12): This field is command specific Dword 12.
55:52	Command Dword 13 (CDW13): This field is command specific Dword 13.
59:56	Command Dword 14 (CDW14): This field is command specific Dword 14.
63:60	Command Dword 15 (CDW15): This field is command specific Dword 15.